arXiv:1412.7834v3 [cs.DS] 30 Mar 2015

Pattern Backtracking Algorithm for the Workflow Satisfiability Problem with User-Independent Constraints

D. Karapetyan[1], A. Gagarin[2], and G. Gutin[2]

[1] University of Nottingham, UK
Daniel.Karapetyan@gmail.com
[2] Royal Holloway, University of London, UK
{Andrei.Gagarin,G.Gutin}@rhul.ac.uk

Abstract. The workflow satisfiability problem (WSP) asks whether there exists an assignment of authorised users to the steps in a workflow specification, subject to certain constraints on the assignment. (Such an assignment is called valid.) The problem is NP-hard even when restricted to the large class of user-independent constraints. Since the number of steps k is relatively small in practice, it is natural to consider a parametrisation of the WSP by k. We propose a new fixed-parameter algorithm to solve the WSP with user-independent constraints. The assignments in our method are partitioned into equivalence classes such that the number of classes is exponential in k only. We show that one can decide, in polynomial time, whether there is a valid assignment in an equivalence class. By exploiting this property, our algorithm reduces the search space to the space of equivalence classes, which it browses within a backtracking framework, hence emerging as an efficient yet relatively simple-to-implement or generalise solution method. We empirically evaluate our algorithm against the state-of-the-art methods and show that it clearly wins the competition on the whole range of our test problems and significantly extends the domain of practically solvable instances of the WSP.

1 Introduction

In the *workflow satisfiability problem* (WSP), we aim at assigning authorised users to the steps in a workflow specification, subject to some constraints arising from business rules and practices. The WSP has applications in information access control (e.g. see [1,2,3]), and it is extensively studied in the security research community [2,3,9,15]. In the WSP, we are given a set U of *users*, a set S of *steps*, a set $\mathcal{A} = \{A(u) : u \in U\}$ of *authorisation lists*, where $A(u) \subseteq S$ denotes the set of steps for which user u is authorised, and a set C of *(workflow) constraints*. In general, a *constraint* $c \in C$ can be described as a pair $c = (T, \Theta)$, where $T \subseteq S$ is the *scope* of the constraint and Θ is a set of functions from T to U which specifies those assignments of steps in T to users in U that satisfy the constraint (authorisations disregarded). Authorisations and constraints described in WSP

literature are relatively simple such that we may assume that all authorisations and constraints can be checked in polynomial time (in $|U|$, $|S|$ and $|C|$).

Given a *workflow* $W = (S, U, \mathcal{A}, C)$, W is *satisfiable* if there exists a function $\pi : S \to U$ such that

- for all $s \in S$, $s \in A(\pi(s))$ (each step is allocated to an authorised user);
- for all $(T, \Theta) \in C$, $\pi|_T \in \Theta$ (every constraint is satisfied).

A function $\pi : S \to U$ is an *authorised* (*eligible*, *valid*, respectively) *complete plan* if it satisfies the first condition above (the second condition, both conditions, respectively).

For example, consider the following instance of WSP. The step and user sets are $S = \{s_1, s_2, s_3, s_4\}$ and $U = \{u_1, u_2, \ldots, u_5\}$. The authorisation lists are $A(u_1) = \{s_1, s_2, s_3, s_4\}$, $A(u_2) = \{s_1\}$, $A(u_3) = \{s_2\}$, $A(u_4) = A(u_5) = \{s_3, s_4\}$. The constraints are $(s_1, s_2, =)$ (the same user must be assigned to s_1 and s_2), $(s_2, s_3, \neq)$ (s_2 and s_3 must be assigned to different users), $(s_3, s_4, \neq)$, and $(s_4, s_1, \neq)$. Since the function π assigning u_1 to s_1 and s_2, u_4 to s_3, and u_5 to s_4 is a valid complete plan, the workflow is satisfiable.

Clearly, not every workflow is satisfiable, and hence it is important to be able to determine whether a workflow is satisfiable or not and, if it is satisfiable, to find a valid complete plan. Unfortunately, the WSP is NP-hard [15] and, since the number k of steps is usually relatively small in practice (usually $k \ll n = |U|$ and we assume, in what follows, that $k < n$), Wang and Li [15] introduced its parameterisation[3] by k. Algorithms for this parameterised problem were also studied in [5,6,7,10]. While in general the WSP is W[1]-hard [15], the WSP restricted[4] to some practically important families of constraints is fixed-parameter tractable (FPT) [6,10,15]. (Recall that a problem parameterised by k is FPT if it can be solved by an FPT algorithm, i.e. an algorithm of running time $O^*(f(k))$, where f is an arbitrary function depending on k only, and O^* suppresses not only constants, but also polynomial factors in k and other parameters of the problem formulation.)

Many business rules are not concerned with the identities of the users that perform a set of steps. Accordingly, we say a constraint $c = (T, \Theta)$ is *user-independent* if, whenever $\theta \in \Theta$ and $\phi : U \to U$ is a permutation, then $\phi \circ \theta \in \Theta$. In other words, given a complete plan π that satisfies c and any permutation $\phi : U \to U$, the plan $\pi' : S \to U$, where $\pi'(s) = \phi(\pi(s))$, also satisfies c. The class of user-independent constraints is general enough in many practical cases; for example, all the constraints defined in the ANSI RBAC standard [1] are user-independent. Most of the constraints studied in [5,7,10,15] and other papers are also user-independent. Classical examples of user-independent constraints are the requirements that two steps are performed by either two different users (*separation-of-duty*), or the same user (*binding-of-duty*). More complex

[3] We use terminology of the recent monograph [11] on parameterised algorithms and complexity.

[4] While we consider special families of constraints, we do not restrict authorisation lists.

constraints state that at least/at most/exactly r users are required to complete some sensitive set of steps (these constraints belong to the family of *counting* constraints), where r is usually small. A simple reduction from Graph Colouring shows that the WSP restricted to the separation-of-duty constraints is already NP-hard [15].

The WSP is an important applied problem and is thoroughly studied in the literature. However, as was shown by Cohen et al. [5], the methods developed so far were capable of solving user-independent WSP instances only for relatively small values of k. In this paper we propose a new approach that, compared to the existing solution methods, significantly extends the number of steps in practically solvable instances now covering the values of k expected in the majority of real-world instances. Importantly, the proposed method is relatively simple to implement or extend with new constraints, such that its accessibility is similar to that of SAT-solvers used by practitioners [15].

The proposed solution method is a deterministic algorithm that uses backtracking to browse the space of all the equivalence classes of partial solutions. We show that it is possible to test efficiently if there exists an authorised plan in a given equivalence class. This makes our algorithm FPT as the number of equivalence classes is exponential in k only.

2 Patterns and the User-Iterative Algorithm

A *plan* is a function $\pi : T \to U$, where $T \subseteq S$ (note that if $T = S$ then π is a complete plan). We define an *equivalence relation* on the set of all plans, which is a special case of an equivalence relation defined in [6]. For user-independent constraints, two plans $\pi : T \to U$ and $\pi' : T' \to U$ are *equivalent*, denoted by $\pi \approx \pi'$, if and only if $T = T'$, and $\pi(s) = \pi(t)$ if and only if $\pi'(s) = \pi'(t)$ for every $s, t \in T$. Assuming an ordering $s_1, s_2, \ldots, s_k$ of steps S, every plan $\pi : T \to U$ can be encoded into a *pattern* $P = P(\pi) = (x_1, \ldots, x_k)$ defined by:

$$x_i = \begin{cases} 0 & \text{if } s_i \notin T, \\ 1 & \text{if } i = 1 \text{ and } s_1 \in T, \\ x_j & \text{if } \pi(s_i) = \pi(s_j) \text{ and } j < i, \\ \max\{x_1, x_2, \ldots, x_{i-1}\} + 1 & \text{otherwise.} \end{cases} \tag{1}$$

The pattern $P(\pi)$ uniquely encodes the equivalence class of π, and $P(\pi) = P(\pi')$ for every π' in that equivalence class [6]. The pattern P represents an assignment of steps in T to some users in any plan of the equivalence class of π. We say that a pattern is *complete* if $x_i \neq 0$ for $i = 1, 2, \ldots, k$.

The state-of-the-art FPT algorithm for the WSP with counting constraints proposed in [5] and called here User-Iterative (UI), iterates over the set of users and gradually computes all encoded equivalence classes of valid plans until it finds a complete solution to the problem, or all the users have been considered. Effectively, it uses the breadth-first search in the space of plans. In the breadth-first search tree, equivalent plans can be generated but they are detected efficiently using patterns and the corresponding search branches are then merged

together. Since the UI algorithm generates a polynomial number of plans per equivalence class and the number of equivalence classes is exponential in k only, the UI algorithm is FPT. The results of [5,7] show that the generic user-iterative FPT algorithm of [6] has a practical value, and its implementations are able to outperform the well-known pseudo-Boolean SAT solver SAT4J [14].

In this paper we propose a new FPT solution method for the WSP which also exploits equivalence classes and patterns but in a more efficient manner. Among other advantages, our algorithm never generates multiple plans within the same equivalence class. For further comparison of our algorithm with the UI algorithm, see Section 4.

3 The Pattern-Backtracking Algorithm

We call our new method *Pattern-Backtracking* (PB) as it uses the backtracking approach to browse the search space of patterns. To describe it, we introduce several additional notations. We will say that a plan $\pi : T \to U$ is *authorised* if $s \in A(\pi(s))$ for every $s \in T$, *eligible* if it does not violate any constraint in C, and *valid* if it is both authorised and eligible. Similarly, a pattern P is *authorised*, *eligible* or *valid* if there exists a plan π such that $P(\pi) = P$ and π is authorised, eligible or valid, respectively. By $P(s_i)$ we denote the value x_i in $P = (x_1, x_2, \ldots, x_k)$. We also use notations $A^{-1}(s) = \{u \in U : \ s \in A(u)\}$ for the set of users authorised for step $s \in S$ and $P^{-1}(x_i) = \{s \in S : \ P(s) = x_i\}$ for all the steps assigned to the same user encoded by the value of x_i in P. Note that $P^{-1}(x_i) \neq \emptyset$ for $i = 1, 2, \ldots, k$ for any complete pattern.

In Section 3.1 we show how to find a valid plan for an eligible pattern, which is an essential part of our algorithm, and in Section 3.2 we describe the algorithm itself.

3.1 Pattern Validity Test

The PB algorithm searches the space of patterns; once an eligible complete pattern P is found, we need to check if it is valid and, if it is, then to find a plan π such that $P = P(\pi)$. The following theorem allows us to address these two questions efficiently.

For a complete pattern $P = (x_1, x_2, \ldots, x_k)$, let $X = \{x_i : \ i = 1, 2, \ldots, k\}$ (note that the cardinality of the set X may be smaller than k). Let $G = (X \cup U, E)$ be a bipartite graph, where $(x_i, u) \in E$ if and only if $u \in A^{-1}(s)$ for each $s \in P^{-1}(x_i)$, $x_i \in X$.

Theorem 1. *A pattern P is authorised if and only if G has a matching of size $|X|$.*

Proof. Suppose M is a matching of size $|X|$ in G. Construct a plan π as follows: for each edge $(x_i, u) \in M$ and $s \in P^{-1}(x_i)$, set $\pi(s) = u$. Since M covers x_i for every $i = 1, 2, \ldots, k$ and P is a complete pattern, the above procedure defines $\pi(s)$ for every step $s \in S$. Hence, π is a complete plan. Now observe that, for each

$x_i \in X$, all the steps $P^{-1}(x_i)$ are assigned to exactly one user, and if $x_i \neq x_j$ for some $i, j \in \{1, 2, \dots k\}$, then $\pi(s_i) \neq \pi(s_j)$ by definition of matching. Therefore $P(\pi) = P$. Observe also that π respects the authorisation lists; for each edge $(x_i, u) \in M \subseteq E$ and each step $s \in P^{-1}(x_i)$, we guarantee that $u \in A(s)$. Thus, plan π is authorised and, hence, pattern $P = P(\pi)$ is also authorised.

On the other hand, assume there exists an authorised plan π such that $P(\pi) = P$ for a given pattern P. Let $X = \{x_i : \ i = 1, 2, \dots, k\}$. Construct a set M as follows: $M = \{(x_i, u) : \ x_i \in X \text{ and } \exists s \in P^{-1}(x_i) \text{ s.t. } u = \pi(s)\}$. Consider a pair $(x_i, u) \in M$, and find some $s \in P^{-1}(x_i)$. Note that, as $P = P(\pi)$ and by definition of pattern, $\pi(s') = u$ for every $s' \in P^{-1}(x_i)$. Since π is authorised, $u \in A(s')$ for every $s' \in P^{-1}(x_i)$, i.e. $(x_i, u) \in E$ and $M \subseteq E$. In other words, M is a subset of edges of G.

Now notice that, for each $x_i \in X$, there exists at most one edge $(x_i, u) \in M$ as $\pi(s') = u$ for every $s' \in P^{-1}(x_i)$. Moreover, for each $u \in U$, there exists at most one edge $(x_i, u) \in M$ as otherwise there would exist some $i, j \in \{1, 2, \dots, k\}$ such that $\pi(s_i) = \pi(s_j)$ and $x_i \neq x_j$, which violates $P = P(\pi)$. Hence, the edge set M is disjoint. Finally, $|M| = |X|$ because $P^{-1}(x_i)$ is non-empty for every $x_i \in X$. We conclude that M is a matching in G of size $|X|$. □

Theorem 1 implies that, to determine whether an eligible pattern P is valid, it is enough to construct the bipartite graph $G = (X \cup U, E)$ and to find a maximum size matching in G. It also provides an algorithm for converting a maximum matching M of size $|X|$ in G into a valid plan π such that $P(\pi) = P$.

The matching problem arising in Theorem 1 has some interesting properties:

- The bipartite graph $G = (X \cup U, E)$ is highly unbalanced as $|X| \leq k$, and we assume that $|U| \gg k$. It is easy to see that the maximum length of an augmenting path in G is $|X| \leq k$ and, hence, the time complexity of the Hungarian and Hopcroft-Karp methods are $O(k^3)$ and $O(k^{2.5})$, respectively.
- We are interested only in matchings of size $|X|$. If the maximum matching is of a smaller size, we do not need to retrieve it.
- Once a matching of size $|X|$ is found, the PB algorithm terminates since a valid plan is found. However, the algorithm might test an exponential number (in k) of graphs with the maximum matching of size smaller than $|X|$. Hence, we are mainly interested in time of checking whether the maximum matching is smaller than $|X|$.

To exploit the above features, we use the Hungarian method with a speed-up heuristic provided by the following proposition.

Proposition 1. *If $M = \{(x_{\chi(1)}, u_{\psi(1)}), \dots, (x_{\chi(t)}, u_{\psi(t)})\}$, $t < |X|$, is a matching in the graph $G = (X \cup U, E)$ such that there exists no M-augmenting path in G starting at a vertex $x_{\chi(t+1)} \in X$, then there is no matching covering all vertices of X in G.*

Proof. W.l.o.g, assume that $M = \{(x_1, u_1), \dots, (x_t, u_t)\}$ and $x_{\chi(t+1)} = x_{t+1}$. Now suppose that G has a matching

$$M' = \{(x_1, y_1), (x_2, y_2), \dots, (x_k, y_k)\}, \ y_i \in U, \ i = 1, 2, \dots, k$$

covering all of X. Consider the symmetric difference of two matchings $H = (M \cup M') \setminus (M \cap M')$. Since every vertex of H has degree at most 2, the graph induced by H consists of some disjoint paths and even cycles having edges alternating between M and M'. Since x_{t+1} is not in M but covered by M', it is an end point of one of the alternating paths in H, say P_{t+1}. Now, it is possible to see that P_{t+1} is an augmenting path in G with respect to the matching M; since, starting at x_{t+1}, every time we use an edge of M' to go from a vertex in X to a vertex in U, P_{t+1} must have an end point in a vertex of U not covered by M (M' covers all the vertices in X). This contradicts the fact that there is no augmenting path in G starting at x_{t+1} with respect to M.□

The result of Proposition 1 allows us to terminate the Hungarian algorithm as soon as a vertex $x_i \in X$ is found such that no augmenting path starting from x_i can be obtained. Construction of the graph takes $O(kn)$ time, and solving the maximum matching problem in G with the Hungarian method takes $O(k^3)$ time.

3.2 The Backtracking Algorithm

The PB algorithm uses a backtracking technique to search the space of patterns, and for each eligible pattern, it verifies whether such a pattern is valid. If a valid pattern P is found, the algorithm returns a complete plan π such that $P(\pi) = P$, see Section 3.1 for details. If no valid pattern is found during the search, the instance is unsatisfiable.

The calling procedure for the PB algorithm is shown in Algorithm 1, which in turn calls the recursive search function in Algorithm 2. The recursive function tries all possible extensions P' of the current pattern P by adding one new step s to it (line 12). The step s is selected heuristically (line 9), where function $\rho(s)$ is an empirically tuned function indicating the importance of step s in narrowing down the search space. The implementation of $\rho(s)$ depends on the specific types of constraints involved in the instance and should reflect the intuition regarding the structure of the problem. See (2) in Section 5 for a particular implementation of $\rho(s)$ for the types of constraints we used in our computational study. Note that our branching heuristic dynamically changes the steps ordering used in the pattern definition in Section 2. Nevertheless, this does not affect any theoretical properties of the pattern.

We use a heuristic (necessary but not sufficient) test (lines 13–15 of Algorithm 2) to check whether the pattern P' can be authorised; that allows us to prune branches which are easily provable to include no authorised patterns.

In line 16, the algorithm checks whether the new pattern P' violates any constraints and, if not, then executes the recursive call.

4 Comparison of the PB and UI Algorithms

In this section we analyse the time and memory complexity of the PB algorithm and compare it to the UI algorithm.

Algorithm 1: Backtracking search initialisation (entry procedure of PB)

input : WSP instance $W = (S, U, \mathcal{A}, C)$
output : Valid plan π or UNSAT
1 Initialise $P(s) \leftarrow 0$ for each $s \in S$;
2 $\pi \leftarrow$ Recursion(P);
3 **return** π (π may be UNSAT here);

Observe that each internal node (corresponding to an incomplete plan) in the search tree of the PB algorithm has at least two children, and each leaf in this tree corresponds to a complete pattern. Thus, the total number of patterns considered by the PB algorithm is less than twice the number of complete patterns. Observe that the number of complete patterns equals the number of partitions of a set of size k, i.e. the kth Bell number B_k. Finally, observe that the PB algorithm spends time polynomial in n on each node of the search tree.[5] Thus, the time complexity of the PB algorithm is $O^*(B_k)$. The PB algorithm follows the depth-first search order and, hence, stores only one pattern at a time. At each leaf node, it also solves the matching problem generating a graph with $O(kn)$ edges. Hence, the memory complexity of the algorithm is $O(kn)$.

It is interesting to compare the PB algorithm to the UI algorithm (briefly described in Section 2). Despite both algorithms using the idea of equivalence classes and being FPT, they have very different working principles and properties.

1. Observe that, in the worst case, the UI algorithm may store all patterns, and the number of patterns is B_{k+1}. Indeed, consider a pattern $P = (x_1, \ldots, x_k)$ and a set $\{s_1, \ldots, s_k, s_{k+1}\}$. Then each partition of the set corresponds to a pattern of P, where $x_i = 0$ if and only if s_i and s_{k+1} are in the same subset of the partition. Therefore, the UI algorithm takes $O(kB_{k+1})$ memory, which is in sharp contrast to the PB algorithm that requires very little memory. Considering that, e.g. $B_{20} = 51\,724\,158\,235\,372$, memory consumption poses a serious bottleneck for the UI algorithm as the RAM capacity of any mainstream machine is well below the value of B_{20}. Moreover, the UI algorithm accesses a large volume of data in a non-sequential order, which might have a dramatic effect on the algorithm's performance when implemented on a real machine as shown in [13].
2. From the practical point of view, the PB algorithm considers less patterns than the UI algorithm ($O(B_k)$ vs. $O(B_{k+1})$) as the PB algorithm assigns the steps in a strict order, avoiding generation of duplicate patterns. Moreover, the PB algorithm generates each pattern at most once, while the UI algorithm is likely to generate a pattern several times rejecting the duplicates afterwards.
3. Both algorithms use heuristics to determine the order in which the search tree is explored. However, while the UI algorithm has to use a certain fixed

[5] Assuming that the WSP instance does not include any exotic constraints.

Algorithm 2: Recursion(P) (recursive function for backtracking search)

input : Pattern P
output : Eligible plan or UNSAT if no eligible plan exists in this branch of the search tree

Initialise the set $S' \subseteq S$ of assigned steps $S' \leftarrow \{s \in S : \ P(s) \neq 0\}$;
if $S' = S$ **then**
    Verify if pattern P is valid (using the matching algorithm of Theorem 1);
    if *pattern P is valid* **then**
        return *plan π realising P*;
    else
        return *UNSAT*;
else
    Select unassigned step $s \in S \setminus S'$ that maximises $\rho(s)$;
    Calculate $k' \leftarrow 1 + \max_{s \in S} P(s)$;
    for $x = 1, 2, \ldots, k'$ **do**
        Set $P(s) \leftarrow x$ to obtain a new pattern P';
        Compute the set of steps Q assigned to x: $Q = \{t \in S : \ P'(t) = x\}$;
        if $|\bigcap_{t \in Q} A^{-1}(t)| = 0$ **then**
            Proceed to the next value of x (reject P');
        if *P' is an eligible pattern* **then**
            $\pi \leftarrow$ Recursion(P');
            if $\pi \neq UNSAT$ **then**
                return π;
return *UNSAT (for a particular branch of recursion; does not mean that the whole instance is unsat)*;

order of users for all the search branches, the PB algorithm has the flexibility of changing the order of steps in each branch of the search. Note that the order of assignments is crucial to the algorithm's performance as it can help to prune branches early.

5 Computational Experiments

In this section we empirically verify the efficiency of the PB algorithm. We compare the following WSP solvers:

PB The algorithm proposed in this paper;
UI Another FPT algorithm proposed in [6] and evaluated in [5,7];
SAT4J A pseudo-Boolean SAT formulation [5,7] of the problem solved with SAT4J.

Due to the difficulty of acquiring real-world WSP instances [5,15], we use the random instance generator described in [5]. Three families of user-independent

constraints are used: *not-equals* (also called *separation-of-duty*) constraints $(s, t, \neq)$, *at-most-r* constraints $(r, Q, \leqslant)$ and *at-least-r* constraints $(r, Q, \geqslant)$. A not-equals constraint $(s, t, \neq)$ is satisfied by a complete plan π if and only if $\pi(s) \neq \pi(t)$. An at-most-r constraint $(r, Q, \leqslant)$ is satisfied if and only if $|\pi(Q)| \leq r$, where Q is the scope of the constraint. Similarly, an at-least-r constraint $(r, Q, \geqslant)$ is satisfied if and only if $|\pi(Q)| \geq r$. We do not explicitly consider the widely used binding-of-duty constraints, that require two steps to be assigned to one user, as those can be trivially eliminated during preprocessing. While the binding-of-duty and separation-of-duty constraints provide the basic modelling capabilities, the at-most-r and at-least-r constraints impose more general "confidentiality" and "diversity" requirements on the workflow, which can be important in some business environments.

The instance generator (available for downloading [12]) takes four parameters: the number of steps k, the number of not-equals constraints e, the number of at-most and at-least constraints c and the random generator seed value. Each instance has $n = 10k$ users. For each user $u \in U$, it generates a uniformly random authorisation list $A(u)$ such that $|A(u)|$ is selected uniformly from $\{1, 2, \ldots, \lceil 0.5k \rceil\}$ at random. It also generates e distinct not-equals, c at-most and c at-least constraints uniformly at random. All at-most and at-least constraints are of the form $(3, Q, \sigma)$, where $|Q| = 5$ and $\sigma \in \{\leqslant, \geqslant\}$.

Our test machine is based on two Intel Xeon CPU E5-2630 v2 (2.6 GHz) and has 32 GB RAM installed. Hyper-threading is enabled, but we never run more than one experiment per physical CPU core concurrently. The PB algorithm is implemented in C#, and the UI algorithm is implemented in C++. Concurrency is not exploited in any of the tested solution methods. The PB algorithm is also available for downloading [12].

The branching heuristic implemented in line 9 of Algorithm 2 selects a step $s \in S$ that maximises a ranking function $\rho(s)$:

$$\rho(s) = c_{\neq}(P) + \alpha c^0_{\leq}(P) + \beta c^1_{\leq}(P) + \gamma c^2_{\leq}(P), \tag{2}$$

where $c_{\neq}(P)$ is the number of not-equals constraints involving step s, $c^i_{\leq}(P)$ is the number of at-most-r constraints involving s such that $r - i$ distinct users are already assigned to it, and α, β and γ are parameters. The intuition is that the steps s that maximise $\rho(s)$ are tightening the search space quickly. The parameters α, β and γ were selected empirically. We found out that the algorithm is not very sensitive to the values of these parameters, and settled down at $\alpha = 100$, $\beta = 2$ and $\gamma = 1$. Note that the function does not account for at-least constraints. This reflects our empirical observation that the at-least constraints are usually relatively weak in our instances and rarely help in pruning branches of search.

We started from establishing what parameter values make the instances hard. However, due to the lack of space, we provide only the conclusions drawn from this series of experiments. As it could be expected, greatly under- and over-subscribed instances are easier to solve than the instances in the region between

those two extremes. The behaviour of the analysed solvers is consistent in this regard. The particular values of the number of not-equals constraints e and the number c of at-most and at-least constraints that make the instances most challenging depend on k. Thus, in our final experiment, which is to establish the maximum size k of instances practically solvable by each of the methods, we considered several instances with a range of parameters to ensure that at least one of them is hard. In particular, we fixed the *density* of not-equals constraints, calculated as $d = \frac{2e}{k(k-1)} \cdot 100\%$, at $d = 10\%$ and the number c of at-most and at-least constraints at each of $c = 1.0k$, $c = 1.2k$ and $c = 1.4k$, producing three instances for each k and seed value.

Each solver is given one hour limitation for each instance from the set. If a solver fails on at least one of the instances (could not terminate within 1 hour), we say that it fails on the whole set. The intention is to make sure that the solver can tackle hard satisfiable and unsatisfiable instances within a reasonable time. The results are presented in Figure 1 in the form of boxplots. The percentage of runs in which the solver succeeded is shown as the width of the box. This information is also provided at the bottom of Figure 1.

The PB algorithm, being faster than the two other methods by several orders of magnitude, reliably solves all the instances of size up to $k = 49$. Compare it to the UI and SAT4J solvers that succeed only for $k \leq 23$ and $k \leq 15$, respectively. Moreover, its running time grows slower than that of the UI and SAT4J solvers, which indicates that it has higher potential if more computational power is allocated. In other words, thanks to our new solution method, the previously unapproachable problem instances of practical sizes can now be routinely tackled.

6 Conclusion

We proposed a new FPT algorithm for the WSP with user-independent constraints. Our experimental analysis have shown that the new algorithm outperforms all the methods in the literature by several orders of magnitude and significantly extends the domain of practically solvable instances. Another advantage of the new FPT algorithm is that it is relatively easy to implement and extend; for example, it is straightforward to parallelise it.

Future research is needed to establish further potential to improve the algorithm's performance. Particular attention has to be paid to the branching heuristic. Thorough empirical analysis has to be conducted to investigate the performance of the algorithms on easy and hard instances.

Another relevant subject was recently studied in [8]; the paper introduces an optimisation version of WSP and proposes an FPT branch and bound algorithm inspired by Pattern Backtracking.

Acknowledgment. This research was partially supported by EPSRC grant EP/K005162/1. The source codes of the Pattern Backtracking algorithm and the instance generator are publicly available [12].

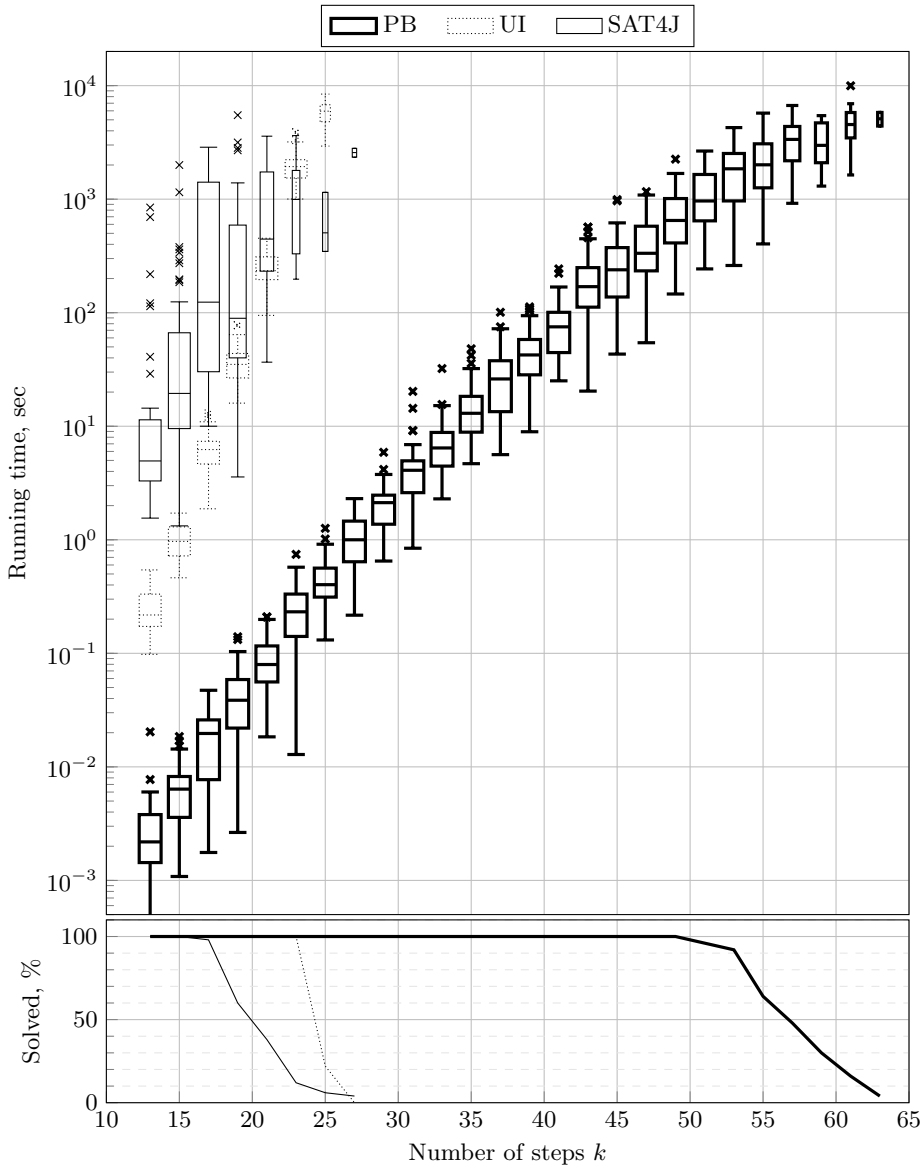

Fig. 1. Running time vs. number of steps k. For each k, we generate 50 instance sets with different seed values. The distributions are presented in the boxplot form, where the width of a box is proportional to the number of instance sets on which the solver succeeded. The plot at the bottom of the figure also shows the success rate of each solver.

References

1. American National Standards Institute. *ANSI INCITS 359-2004 for Role Based Access Control*, 2004.
2. Basin, D.A., Burri, S.J., Karjoth, G.: Obstruction-free authorisation enforcement: Aligning security and business objectives. Journal of Computer Security 22(5), 661–698 (2014).
3. Bertino, E., Ferrari, E., Atluri, V.: The specification and enforcement of authorisation constraints in workflow management systems. ACM Trans. Inf. Syst. Secur. 2(1), 65–104 (1999)
4. Chimani, M., Klein, K.: Algorithm engineering: Concepts and practice. In: Bartz-Beielstein, T., Chiarandini, M., Paquete, L., Preuss, M. (eds.) Experimental methods for the analysis of optimization algorithms, pp. 131–158 (2010)
5. Cohen, D., Crampton, J., Gagarin, A., Gutin, G., Jones, M.: Engineering algorithms for workflow satisfiability problem with user-independent constraints. Proc. 8th International Frontiers of Algorithmics Workshop (FAW 2014), J. Chen, J.E. Hopcroft, and J. Wang (Eds.), 2014, LNCS 8497, Springer, pp. 48–59.
6. Cohen, D., Crampton, J., Gagarin, A., Gutin, G., Jones, M.: Iterative plan construction for the workflow satisfiability problem. *Journal of Artificial Intelligence Research* 51, pp. 555–577 (2014)
7. Cohen, D., Crampton, J., Gagarin, A., Gutin, G., Jones, M.: Algorithms for the workflow satisfiability problem engineered for counting constraints, to appear in J. of Combin. Optim., (2014)
8. Crampton, J., Gutin, G., Karapetyan, D.: Valued Workflow Satisfiability Problem, to appear in proc. of ACM symposium on Access control models and technologies (SACMAT), 13 June, Vienna, Austria. ACM (2015)
9. Crampton, J.: A reference monitor for workflow systems with constrained task execution. In: Ferrari, E., Ahn, G.J. (eds.) SACMAT. pp. 38–47. ACM (2005)
10. Crampton, J., Gutin, G., Yeo, A.: On the parameterized complexity and kernelization of the workflow satisfiability problem. ACM Trans. Inf. Syst. Secur. 16(1), 4 (2013)
11. Downey, R.G., Fellows, M.R.: Foundations of Parameterized Complexity. Springer (2013)
12. Karapetyan, D., Gutin, G., Gagarin, A.: Source codes of the Pattern Backtracking algorithm and the instance generator. DOI:`10.6084/m9.figshare.1360237` retrieved 31 March 2015.
13. Karapetyan, D., Gutin, G., Goldengorin, B.: Empirical evaluation of construction heuristics for the multidimensional assignment problem. Proc. London Algorithmics 2008: Theory and Practice, 2009, Texts in Algorithmics 11, College Publications, pp. 107–122.
14. Le Berre, D., Parrain, A.: The SAT4J library, release 2.2. J. Satisf. Bool. Model. Comput. 7, 59–64 (2010)
15. Wang, Q., Li, N.: Satisfiability and resiliency in workflow authorisation systems. ACM Trans. Inf. Syst. Secur. 13(4), 40 (2010)